

McGINN & GIBB, PLLC
A PROFESSIONAL LIMITED LIABILITY COMPANY
PATENTS, TRADEMARKS, COPYRIGHTS, AND INTELLECTUAL PROPERTY LAW
8321 OLD COURTHOUSE ROAD, SUITE 200
VIENNA, VIRGINIA 22182-3817
TELEPHONE (703) 761-4100
FACSIMILE (703) 761-2375; (703) 761-2376

**APPLICATION
FOR
UNITED STATES
LETTERS PATENT**

APPLICANT: **GUSTAVSON, ET AL.**

FOR: **METHOD AND STRUCTURE FOR
PRODUCING HIGH PERFORMANCE
LINEAR ALGEBRA ROUTINES USING
STREAMING**

DOCKET NO.: **YOR920030331US1**

METHOD AND STRUCTURE FOR PRODUCING HIGH PERFORMANCE LINEAR ALGEBRA ROUTINES USING STREAMING

5

DESCRIPTION

BACKGROUND OF THE INVENTION

Cross-Reference to Related Applications

The following seven Applications, including the present Application, are related:

10

1. U.S. Patent Application No. 10/ __, __, filed on ____, to Gustavson et al., entitled "METHOD AND STRUCTURE FOR PRODUCING HIGH PERFORMANCE LINEAR ALGEBRA ROUTINES USING COMPOSITE BLOCKING BASED ON L1 CACHE SIZE", having IBM Docket YOR920030010US1;

15

2. U.S. Patent Application No. 10/ __, __, filed on ____, to Gustavson et al., entitled "METHOD AND STRUCTURE FOR PRODUCING HIGH PERFORMANCE LINEAR ALGEBRA ROUTINES USING A HYBRID FULL PACKED STORAGE FORMAT", having IBM Docket YOR920030168US1;

20

3. U.S. Patent Application No. 10/ __, __, filed on ____, to Gustavson et al., entitled "METHOD AND STRUCTURE FOR PRODUCING YOR920030331US1

HIGH PERFORMANCE LINEAR ALGEBRA ROUTINES USING REGISTER
BLOCK DATA FORMAT”, having IBM Docket YOR920030169US1;

4. U.S. Patent Application No. 10/ __, __, filed on __, __, to
Gustavson et al., entitled “METHOD AND STRUCTURE FOR PRODUCING
5 HIGH PERFORMANCE LINEAR ALGEBRA ROUTINES USING LEVEL 3
PREFETCHING FOR KERNEL ROUTINES”, having IBM Docket
YOR920030170US1;

5. U.S. Patent Application No. 10/ __, __, filed on __, __, to
Gustavson et al., entitled “METHOD AND STRUCTURE FOR PRODUCING
10 HIGH PERFORMANCE LINEAR ALGEBRA ROUTINES USING
PRELOADING OF FLOATING POINT REGISTERS”, having IBM Docket
YOR920030171US1;

6. U.S. Patent Application No. 10/ __, __, filed on __, __, to
Gustavson et al., entitled “METHOD AND STRUCTURE FOR PRODUCING
15 HIGH PERFORMANCE LINEAR ALGEBRA ROUTINES USING A
SELECTABLE ONE OF SIX POSSIBLE LEVEL 3 L1 KERNEL ROUTINES”,
having IBM Docket YOR920030330US1; and

7. U.S. Patent Application No. 10/ __, __, filed on __, __, to
Gustavson et al., entitled “METHOD AND STRUCTURE FOR PRODUCING
20 HIGH PERFORMANCE LINEAR ALGEBRA ROUTINES USING

YOR920030331US1

STREAMING”, having IBM Docket YOR920030331US1, all assigned to the present assignee, and all incorporated herein by reference.

Field of the Invention

5 The present invention relates generally to a technique for improving performance for linear algebra routines, with special significance to optimizing the matrix multiplication process as exemplarily implemented in the existing LAPACK (Linear Algebra PACKage) standard. More specifically, a streaming technique allows a steady and timely flow of matrix data from different cache levels, in which submatrices of A, B, and C "play the role" of scalar, vector, and 10 matrix in a general linear algebra subroutine kernel and are selectively stored in different cache levels.

Description of the Related Art

Scientific computing relies heavily on linear algebra. In fact, the whole field of engineering and scientific computing takes advantage of linear algebra for 15 computations. Linear algebra routines are also used in games and graphics rendering.

Typically, these linear algebra routines reside in a math library of a computer system that utilizes one or more linear algebra routines as a part of its processing. Linear algebra is also heavily used in analytic methods that include

YOR920030331US1

applications such as supply chain management, as well as numeric data mining and economic methods and models.

A number of methods have been used to improve performance from new or existing computer architectures for linear algebra routines.

5 However, because linear algebra permeates so many calculations and applications, a need continues to exist to optimize performance of matrix processing.

SUMMARY OF THE INVENTION

10 In view of the foregoing and other exemplary problems, drawbacks, and disadvantages, it is, therefore, an exemplary feature of the present invention to provide various techniques that improve performance for linear algebra routines.

 It is another exemplary feature of the present invention to improve factorization routines which are key procedures of linear algebra matrix
15 processing.

 It is another exemplary feature of the present invention to provide more efficient techniques to access data in linear algebra routines.

 It is another exemplary feature of the present invention to provide a method of streaming data for matrix multiplication routines in which matrices are
20 placed into caches based on a size of the matrices.

YOR920030331US1

To achieve the above and other exemplary purposes, in a first exemplary aspect of the present invention, described herein is a method (and structure) for executing a linear algebra subroutine on a computer having a cache (e.g., an L1 cache), including streaming data for matrices involved in processing the linear algebra subroutine such that data is processed using data for a first matrix stored in the L1 cache as a matrix format and data from a second matrix and a third matrix is stored in a memory device at a higher level than the L1 cache, the streaming providing data from the higher level as the streaming data is required for the processing.

In a second exemplary aspect of the present invention, also described herein is a signal-bearing medium tangibly embodying a program of machine-readable instructions executable by a digital processing apparatus to perform the above-described method.

In a third exemplary aspect of the present invention, also described herein is a method of at least one of solving and applying a scientific/engineering problem, including at least one of: using a linear algebra software package that performs one or more matrix processing operations, the method comprising streaming data for matrices involved in processing the linear algebra subroutines such that data is processed using data for a first matrix stored in a cache (e.g., an L1 cache) as a matrix format and data from a second matrix and a third matrix is stored in a memory device at a higher level than the L1 cache, the streaming

YOR920030331US1

providing data from the higher level in a manner as the streaming data is required for the processing; providing a consultation for purpose of solving a scientific/engineering problem using the linear algebra software package; transmitting a result of the linear algebra software package on at least one of a network, a signal-bearing medium containing machine-readable data representing the result, and a printed version representing the result; and receiving a result of the linear algebra software package on at least one of a network, a signal-bearing medium containing machine-readable data representing the result, and a printed version representing the result.

The present invention provides a technique for efficiently bringing data for matrix processing through multiple levels of a computer memory system. By using a generalized streaming technique, data for one matrix in a Level 3 linear algebra subroutine will reside in a first level of cache, and data for the two remaining matrices will be streamed into the L1 cache as required by the subroutine directing the processing.

BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing and other exemplary purposes, aspects and advantages will be better understood from the following detailed description of exemplary embodiments of the invention with reference to the drawings, in which:

YOR920030331US1

Figure 1 illustrates a matrix representation 100 for an operation exemplarily discussed herein;

Figure 2 illustrates an exemplary hardware/information handling system 200 for incorporating the present invention therein;

5 Figure 3 illustrates an exemplary CPU/Floating Point Unit (FPU) architecture 302 which can be incorporated into the present invention;

Figure 4 exemplarily illustrates in more detail the CPU 211 which can be used in a computer system 200 for the present invention;

10 Figure 5 exemplarily illustrates a picture 500 of the matrix layout for L2 and L1 for an example discussed herein;

Figure 6 exemplarily illustrates a larger picture 600 for L3 and L2 for the same example discussed herein; and

Figure 7 illustrates a signal bearing medium 700 (e.g., storage medium) for storing steps of a program of a method according to the present invention.

15

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS OF THE INVENTION

Referring now to the drawings, and more particularly to Figure 1, a preferred embodiment of the present invention will now be discussed. The present invention addresses efficiency in the calculations of linear algebra

20

YOR920030331US1

5 routines, as shown exemplarily in Figure 1, in which is illustrated a processing of the matrix operation $C = C - A^T * B$. For purpose of discussion only, Level 3 BLAS (Basic Linear Algebra Subprograms) of the LAPACK (Linear Algebra PACKage) are used, but it is intended to be understood that the concepts discussed herein are easily extended to other linear algebra mathematical operations and math library modules.

10 Figure 1 illustrates processing of an exemplary matrix operation 100 (e.g., $C = C - A^T * B$). In processing this operation, matrix A is first transposed to form transpose-matrix-A (e.g., A^T) 101. Next, transposed matrix A^T is multiplied with matrix B 102 and then subtracted from matrix C 103. The computer program executing this matrix operation will achieve this operation using three loops 104 in which the element indices of the three matrices A, B, C will be varied in accordance with the desired operation.

15 That is, as shown in the lower section of Figure 1, the inner loop and one step of the middle loop will cause indices to vary so that MB rows 105 of matrix A^T will multiply with NB columns 106 of matrix B. The index of the outer loop will cause the result of the register block row/column multiplications to then be subtracted from the MB-by-NB submatrix 107 of C to form the new submatrix 107 of C. Figure 1 shows an exemplary "snapshot" during execution of one step of the middle loop $i = i: i + MB - 1$ and all steps of the inner loop l, with the outer loop $j = j: j + NB - 1$.

20 YOR920030331US1

A key idea of the generalized streaming discussed herein is the generalization of Example 1 shown in Figure 1 for higher levels of the memory hierarchy. In Example 1, there is a cache resident matrix A^T of size $K1 \times M1$, a vector block B of size $K1 \times N0$, and a scalar block C of size $M0 \times N0$. At a given instant of time, the sizes of the vector block and scalar block is of order 0 and the size of the A is of order 1 (see bottom of Figure 5).

All of a piece of C and part of a piece of B has size of order 2, and these pieces of C and B reside in the next higher level of the memory hierarchy, namely L2 cache. In L2 cache we have cache resident matrix C of size $M2 \times N2$ and, at a given instant in time, streaming pieces (vector block B of size $K1 \times N2$) and scalar block A of size $M1 \times K1$ (e.g., see top of Figure 5).

Before stopping here, we mention briefly TLB (Translation Lookaside Buffer). The TLB is a term well understood in the art as referring to accessing data residing in the memory spanned by the current TLB pages. Usually, the size of the TLB is a sizable fraction of the size of the L2 cache.

Thus, in the above paragraph, the size of $N2$ is chosen so that the cache resident piece and the two streaming pieces can "fit" into the TLB. By doing so, TLB thrashing can be avoided.

Returning, the generalization should be clear. In level 3 cache, we have cache resident matrix piece of A of size $M3 \times K3$ and, at a given instance of time,

YOR920030331US1

a vector block piece of B of size $K3 \times N2$ and a scalar block piece of C of size $M2 \times N2$ (e.g., see Figure 6, and note that is exactly like the bottom of Figure 5).

However, before presenting the exemplary details of the present invention, the following general discussion provides a background of linear algebra subroutines and computer architecture as related to the terminology used herein.

Linear Algebra Subroutines

The explanation of the present invention includes reference to the computing standard called LAPACK (Linear Algebra PACKage) and to various subroutines contained therein. LAPACK is well known in the art and information on LAPACK is readily available on the Internet. When LAPACK is executed, the Basic Linear Algebra Subprograms (BLAS), unique for each computer architecture and provided by the computer vendor, are invoked. LAPACK comprises a number of factorization algorithms for linear algebra processing.

For example, Dense Linear Algebra Factorization Algorithms (DLAFAs) include matrix multiply subroutine calls, such as Double-precision Generalized Matrix Multiply (DGEMM). At the core of level 3 Basic Linear Algebra Subprograms (BLAS) are “L1 kernel” routines which are constructed to operate at near the peak rate of the machine when all data operands are streamed through or reside in the L1 cache.

YOR920030331US1

The most heavily used type of level 3 L1 DGEMM kernel is Double-precision A Transpose multiplied by B (DATB), that is, $C = C - A^T * B$, where A, B, and C are generic matrices or submatrices, and the symbology A^T means the transpose of matrix A (see Figure 1). It is noted that DATB is commonly the only such kernel employed by today's state of the art codes, although DATB is only one of six possible kernels.

The DATB kernel operates so as to keep the A operand matrix or submatrix resident in the L1 cache. Since A is transposed in this kernel, its dimensions are K1 by M1, where K1 x M1 is roughly equal to the size of the L1. Matrix A can be viewed as being stored by row, since in Fortran, a non-transposed matrix is stored in column-major order and a transposed matrix is equivalent to a matrix stored in row-major order. Because of asymmetry (C is both read and written) K1 is usually made to be greater than M1, as this choice leads to superior performance.

Exemplary Computer Architecture

Figure 2 shows a typical hardware configuration of an information handling/ computer system 200 usable with the present invention. Computer system 200 preferably has at least one processor or central processing unit (CPU) 211. Any number of variations are possible for computer system 200, including

YOR920030331US1

various parallel processing architectures and architectures that incorporate one or more FPUs (floating-point units).

In the exemplary architecture of Figure 2, the CPUs 211 are interconnected via a system bus 212 to a random access memory (RAM) 214, read-only memory (ROM) 216, input/output (I/O) adapter 218 (for connecting peripheral devices such as disk units 221 and tape drives 240 to the bus 212), user interface adapter 222 (for connecting a keyboard 224, mouse 226, speaker 228, microphone 232, and/or other user interface device to the bus 212), a communication adapter 234 for connecting an information handling system to a data processing network, the Internet, an Intranet, a personal area network (PAN), etc., and a display adapter 236 for connecting the bus 212 to a display device 238 and/or printer 239 (e.g., a digital printer or the like).

Although not specifically shown in Figure 2, the CPU of the exemplary computer system could typically also include one or more floating-point units (FPU's) that performs floating-point calculations. Computers equipped with an FPU perform certain types of applications much faster than computers that lack one. For example, graphics applications are much faster with an FPU. An FPU might be a part of a CPU or might be located on a separate chip. Typical operations are floating point arithmetic, such as fused multiply/add (FMA), addition, subtraction, multiplication, division, square roots, etc.

YOR920030331US1

Details of the FPU are not so important for an understanding of the present invention, since a number of configurations are well known in the art. Figure 3 shows an exemplary typical CPU 211 that includes at least one FPU 301. The FPU function of CPU 211 includes controlling the FMAs (floating-point multiply/add) and at least one load/store unit (LSU) 301 to service a number of floating point registers (FRegs) 303.

It is noted that in the context of the present invention involving linear algebra processing, the term "FMA" can also be translated as "fused multiply-add" operation/unit, and the present discussion does not care which of the translations is used. The role of the LSU 301 is to move data from a memory device 304 to the registers 303 used by the FMA and to subsequently transfer the results back into a memory device 304.

Another important exemplary aspect of the present invention relates to computer architecture that incorporates a memory hierarchy involving one or more cache memories. Figure 4 shows in more detail how the simple computer system 200 might incorporate a cache 401 in the CPU 211.

Discussion of the present invention includes reference to levels of cache, and more specifically, level 1 cache (L1 cache), level 2 cache (L2 cache) and even level 3 cache (L3 cache). Level 1 cache is typically considered as being a cache that is closest to the CPU, and might even be included as a component of the

YOR920030331US1

CPU, as shown in Figure 4. A level 2 (and higher-level) cache is typically considered as being cache outside the CPU.

5 The details of the cache structure and the precise location of the cache levels are not so important to the present invention so much as recognizing that memory is hierarchical in nature in modern computer architectures and that matrix computation can be enhanced considerably by modifying the processing of matrix subroutines to include considerations of the memory hierarchy.

10 Additionally, in the present invention, it is preferable that the matrix data be laid out contiguously in memory in "stride one" form. "Stride one" means that the data is preferably contiguously arranged in memory to honor double-word boundaries and the useable data is retrieved in increments of the line size.

15 The importance of having six kernel types available is that stride one memory access is desirable for matrix processing. The matrices A and B are usually stored either by row or by column. By having six kernel types, one can choose a kernel in which stride one is available for both operands. Having only one kernel type, instead of six kernel types, means that data copy must be done to provide the format of the one conventional kernel. This means a certain performance loss that might have to be repeated several times.

YOR920030331US1

Streaming Level 3 L1 Kernel Routines for General Matrix Multiplication

All modern programming languages (C, Fortran, etc.) store matrices as two-dimensional arrays for a matrix A. That is, let matrix A have M rows and N columns. The standard column major format of A is as follows.

5 Each of the N columns of A are stored as a contiguous vector (stride one).
Each of the M rows of A are stored with consecutive elements separated by LDA
(Leading Dimension of A) storage locations (stride LDA). Let A(0,0) be stored in
memory location α . The matrix element A(i,j) is stored in memory location $\alpha + i$
+ LDA*j. It is important to note here that stride one is optimal for memory
10 accesses and that stride LDA is poor. Also, almost all level 3 linear algebra code
treats rows and columns about equally.

Modern architectures support this programming language construct.
Further, existing codes support this data structure. Thus, the majority of
programming architectures support, in their machine (engineering)
15 implementations, machine (optimized) instructions that can be used to easily
address and, hence, quickly access elements of one- and two-dimensional arrays
that are stored in traditional formats. At the register level, the register blocks are
submatrices stored in a fashion to optimally exploit the FPU enhanced hardware.

The underlying program construct that supports one- and two-dimensional
20 arrays is the FOR or DO loop. This is an example of generalized counting via
arithmetic progression:

YOR920030331US1


```

DO i = istart, iend, istep
...
ENDDO.

```

5 When a contiguous array that fits in a cache (e.g., the L1 cache) is accessed using this programming construct, near-peak performance can be obtained (provided there is a good L1 cache-FPU register interface) because the architectures are optimized to be as efficient as possible.

10 Traditional IJL (FOR) loop implementations of level 3 L1 DGEMM are likely to be optimal when the underlying arrays are contiguous and map well into the L1 cache. Thus, there is at most six ($3! = 3 \times 2$) level 3 L1 DGEMM kernels to consider. Despite conventional practice, which employs only one, all six are viable under appropriate circumstances. Stride-one data access is the most efficient form of accessing data. Application of the fact is less important when the entire two-dimensional arrays reside in L1 cache as, often, for data in the L1

15 cache, access time behaves in a manner consistent with (uniform) RAM (Random Access Memory).

Thus, in the context of the present invention, "streaming" occurs when parts of matrices fit in the L2 cache, rather than the L1 cache, and will be accessed systematically as explained below. In other words, typically only one of the

20 matrices, or a submatrix thereof, is L1 cache-resident .

A pattern of all six kernel types is to have an L1 cache resident matrix operand, while the remaining two matrix operands (called "streaming matrices")

YOR920030331US1

reside in the next higher memory level (e.g., L2 in this case), if necessary. The outer loop variable defines the common order of streaming. The size of the common dimension of the two streaming matrices is determined by the size of L2.

That is, if N is assumed to be the streaming dimension, the two streaming matrices are C and B , whose dimensions are M by N and K by N . Storage for C and B is MN and KN , respectively. If cache L2 has size $L2S$, then, assuming that C is the cache resident matrix of L2, $MN < L2S$ or $N < L2S/M$.

The discussion in the sixth of the above-identified co-pending Applications provided a brief discussion on streaming for the six kernel types. At level L1, there is a cache resident matrix and two matrices streaming from level L2 (see also Figure 1).

This discussion herein continues the discussion of streaming by describing the concept of generalized streaming and, more specifically, how to choose the cache resident matrix at level L2 (more generally, at level $i + 1$) from the remaining two choices after the matrix has been selected for L1 (more generally, level i).

As an example, the IJL forms of DGEMM leading to DAXPY (scaled vector addition) and DDOT (dot-product) forms are now presented. For example, the equation $C = C - A * B$ can be considered with either:

DDOT $\rightarrow C = C - A^T * B$ (e.g., transpose A and use DATB)

DAXPY $\rightarrow C = C - A * B$ (e.g., copy A and use DAB kernel)

YOR920030331US1

Any IJL form leads to a cache residency pattern of matrices A, B, and C of type matrix, vector, and scalar (e.g., more specifically, patterns where submatrices of A, B, and C "play the role" of scalar, vector, and matrix in the kernel type).

Thus, to be more specific, in the context of the present invention, one is using streaming where the vector and scalar parts of matrices fit in the L2 (not L1) cache. In other words, only the submatrix playing the role of the matrix is L1 cache-resident.

To illustrate the technique of streaming, one of the DDOT kernels is used. However, the same holds true for the other DDOT kernel and all 4 DAXPY kernels. A simplified DATB kernel (JIL order) (see also Figure 1), including 4-by-4-by-1 register blocking, follows:

```

do j = 0, N-4, 4
  do i = 0, M-4, 4
    Tij = C(i:i+3, j:j+3)
    do l = 0, K-1
      Ai = A(l, i:i+3)
      Bj = A(l, j:j+3)
      Tij = Tij + Ai * Bj
    enddo
    C(i:i+3, j:j+3) = Tij
  enddo
enddo

```

! 16 loads of C
! 4 loads of A^T
! 4 loads of B
! 16 FMAs
! 16 stores of C

In the DATB kernel illustrated above, consider the outer *do* loop on *j* as removed and then consider one iteration (*i* = *i* -> *i*+3) of the middle loop. The matrix data necessary for this iteration is A(0:K-1, *i*:*i*+3), B(0:K-1, *j*:*j*+3), and

YOR920030331US1

$C(i:i+3,j:j+3)$. During this execution time interval, 16 elements of C are needed, as are four vectors, $B(0:K-1,j:j+3)$. (see also Figure 1.)

Now, consider the entire middle loop with J fixed at $j:j+3$. During this larger execution time interval, all of $A(0:K-1,0:M-1)$ matrix, a 4 vector of $B(0:K-1,j:j+3)$, and scalar pieces ($C(i:i+3,j:j+3)$, read in and written out) of $C(0:M-1,j:j+3)$ are required.

Thus, during one iteration of the outer loop, cache residency includes a scalar piece or pieces of C , a vector piece of B , and the entire contiguous matrix A . It follows that the entire execution of DATB (over the outer loop) includes streaming vectors of B and scalars of C from $L2$ into and out of $L1$, during which the entire matrix A remains cache resident.

It is noted that tiny pieces of B and C that reside in $L1$ during one iteration of the outer loop are reused M and K times, respectively. Thus, their initial loading penalty from $L2$ is amortized by those selfsame factors.

The remaining five loop permutations: IJL , ILJ , LIJ , JLI , and LJI each give rise to matrix, vector, and scalar patterns of input matrices BAC , BCA , CBA , ACB , and CAB , respectively. The last four patterns are DAXPY kernel types while the first pattern is the second DDOT kernel type.

As mentioned above, the pattern of all six kernel types is to have an $L1$ cache resident matrix operand while the remaining two matrix operands (called "streaming matrices") reside, either entirely or in pieces, in the next higher

YOR920030331US1

memory level (e.g., L2 in this case). The outer loop variable, N for the DDOT kernel above, defines the common order of streaming. The size of the common dimension of the two streaming matrices (e.g., B and C has common dimension N in the DDOT example above) is determined by the size of L2 (e.g., $N < L2S/M$ if C is the resident matrix of L2).

It follows that, with data copy, which is a one-time cost, these IJL ideas generalize to all $l = 1, \dots, M$ caches, plus main memory (i.e., a total of $M+1$ levels). For purpose of the present invention, this more general concept of streaming is referred to as "generalized streaming", as further discussed below.

10 **Generalized Streaming/Optimal multi-level blocking**

In the discussion above, streaming is described for six types of kernels. At level i , $i = 1$ in the first embodiment, there is a cache-resident matrix and two streaming matrices from level $i+1$. The following discussion describes how to choose the cache resident matrix at level $i+1$ from the two possible choices for $i > 1$.

First, matrices A, B, and C have orders M by K, K by N, and M by N for $C = C+A*B$. There are six possible arrangements of the ordering M, N, and K, assuming $M \neq N$, $K \neq N$, $M \neq K$:

20	$M > N > K$	$N > M > K$	$K > M > N$
	$M > K > N$	$N > K > M$	$K > N > M$

YOR920030331US1

Streaming occurs optimally only if the common dimension can be large. The reason for this is conservation of the matrix data. Hence for level one, the largest order is chosen as the streaming dimension, for example, $a = N2$ (in general, $a = M$ or N or K). Here, the suffix "2" on N indicates that the cache resident matrix (i.e. the matrix that remains in that level of the cache) should occupy most of the L2 cache. It is noted that, if a (N in this case) is small enough, then up to L2 level blocking will suffice.

Moving now to level 2, when a is sufficiently large, there are two streaming matrices with common dimension a . The remaining orders are β and δ for the dimensions of the two other streaming matrix dimensions. From the two remaining matrices, the one which is chosen has other dimension $\gamma = \max(\beta, \delta)$.

Hence, streaming at level 2 will be γ , the second largest order (from $M/N/K$). This defines a "generalized" kernel at level 2 (one of the six choices) described above. The L2 cache resident matrix has size a ($N2$) by η ($M2$) with streaming dimension γ ($K3$).

In this example, $\eta = M2$ is the remaining order from M , N , and K . In L1 the cache resident matrix has order η ($M1$) by γ ($K1$) with streaming dimension a ($N2$). The size of L3 determines how much streaming can be allocated to the L2 streaming dimension γ (K in this case) whenever γ is sufficiently large. Of course, if γ is small enough then blocking up to L3 will prove sufficient.

YOR920030331US1

For illustration, it is assumed that $N > K > M$. The L1 kernel will be JIL (L is the index for the K dimension). See also the example in Figure 1. The cache resident matrix, A, will have size m_1 by k_1 (and it is stored in transposed format to facilitate stride one access) and streaming matrix pieces of size m_1 by n_0 (C) and k_1 by n_0 (B) will enter from L2. Note that m_1 by n_0 (C) enters in smaller pieces m_0 by n_0 . The L2 cache resident matrix will be m_2 by n_2 C along with streaming matrix pieces of size m_2 by k_1 A (scalar) and k_1 by n_2 B (vector). Again, note that m_2 by k_1 (A) is made up of scalar pieces of size m_1 by k_1 . The inner loop of the generalized kernel will be a call to the kernel routine, JIL which has loop variables j_1, i_1, l_1 .

Hence, denoting loop variables l_2, i_2, j_2 for the generalized L2 kernel, it can be seen that the j_2 loop of the generalized L2 kernel collapses with the j_1 loop variable of the L1 kernel JIL under this new type of L2-L1 blocking for matrix multiply. We have loop variables l_2, i_2, j_1, i_1, l_1 or l_2, i_2 (L1 kernel routine call).

In Figure 5, an exemplary picture 500 is given of the matrix layout for L2 and L1, and the pseudocode listing below gives the matrix multiply routine that would correspond to the layout depicted in Figure 5. It is noted again that the numerical suffix is associated with a cache level (e.g., K1, M2). The outermost l_2 loop of the pseudocode below indicates that $M2 \times N2$ C 501, $K1 \times N2$ B 502, and $M2 \times K1$ A 503 will occupy L2 and that during this execution time interval $\text{Ceil}[K/K1] * \text{Ceil}[M/M1] * \text{Ceil}[N/N2]$ calls will be made to kernel DATB,

YOR920030331US1

where for each call, L1 cache residency will be $M1 \times K1$ A 504, $K1 \times N0$ B 505, and $M1 \times N0$ C 506. ($M1 \times N0$ C is made up of scalar $M0 \times N0$ pieces of C.)

Pseudocode for Figure 5

```

5      DO l2 = 0, K-1, k1
      lb2 = min(K-l2,k1)
      DO i2 = 0, M-1, m1
      ib2 = min(M-i2, m1)
      DO j1 = 0, N-1, n2
      jb1 = min(N-j1,n2)
10      call datb(ib2,jb1,lb2,A(l2,i2),lda,B(l2,j1),ldb,C(i2,j1),ldc)
      ENDDO
      ENDDO
      ENDDO

```

This L2-L1 description generalizes to an L+1 level (L caches and main
15 memory) memory hierarchy. In the above example, the L3 picture would be a
(larger) replica 600 of the L1 picture, as shown exemplarily in Figure 6. This
should also be compared with Figure 1 for the smaller picture. L3 residency is
made up of $M3 \times K3$ matrix A 601, $K3 \times N2$ vector B 602, and " $M3 \times N2$ " scalar C
603.

20 It is noted that $M1$, $N0$, $K1$ get replaced by $M3$, $N2$, and $K3$, respectively,
as this JIL kernel in L3 is a replica of the same JIL kernel in L1, two levels below
L3. Likewise, the L4 picture would be a replica of the L2 picture where $M2$, $N2$,
and $K3$ are replaced by $M4$, $N4$, and $K5$, respectively. Here, the replica "drop" is
from L4 to L2.

YOR920030331US1

So, in an $L+1$ level memory hierarchy, we would have two generalized kernel routines that flip flop result. The multi-level blocking routine would have two loop variables that block the cache-resident matrix at level l ($l = 1, \dots, L+1$). Each kernel has three loop variables. However, due to the nature of streaming,
5 loop collapsing is obtained between every level-pair.

In the current example, $N > K > M$, JIL is obtained at levels 1,3,5,... and LIJ at levels 2,4,6,... Between levels $2i$ and $2i-1$, one of the two J loop variables disappears. Between levels $2i+1$ and $2i$, one of the two L loop variables disappears.

10 Previous work reduced the poly-algorithmic count from 6^m to 2^m , where $m = M + 1$ is the number of memory hierarchy levels. Furthermore, previous work emphasized only one of six kernels but stating that a second DDOT kernel type was possible. Here, it has been shown that this complexity can be reduced to six, irrespective of the number of levels in the memory hierarchy. Also, at the L1
15 level, this invention admits (see also the sixth above-listed copending application) that six kernel types are possible or desirable instead of just one or two kernel types. Simply select algorithms X and Y (as optimal for the lowest two levels of the hierarchy, based on the input matrix sizes). Then switch back and forth (flip-flop) between algorithms X and Y as the memory hierarchy is traversed
20 upward.

YOR920030331US1

It should be apparent that, in general, in an exemplary preferred embodiment, the computer will have a plurality i of caches, $i = 1, \dots, M+1$. Data for matrices involved in processing a linear algebra subroutine will be streamed such that data for a first matrix will be resident in level i cache, either in its entirety or as a submatrix, and data from a second matrix and third matrix will be stored in the higher level $i+1$ cache, also either in the entirety or as submatrices, and the streaming will provide data as required for the processing.

It should also be apparent that the concepts described herein can readily be combined with one or more of the concepts taught in the remaining co-pending Applications to achieve even more efficiency than when incorporated as an isolated technique.

Software Product Embodiments

In addition to the hardware/software environment described above, a different exemplary aspect of the invention includes a computer-implemented method for performing the invention.

Such a method may be implemented, for example, by operating a computer, as embodied by a digital data processing apparatus, to execute a sequence of machine-readable instructions. These instructions may reside in various types of signal-bearing media.

Thus, this exemplary aspect of the present invention is directed to a programmed product, comprising signal-bearing media tangibly embodying a

YOR920030331US1

program of machine-readable instructions executable by a digital data processor incorporating the CPU 211 and hardware above, to perform the method of the invention.

5 This signal-bearing media may include, for example, a RAM contained within the CPU 211, as represented by the fast-access storage for example. Alternatively, the instructions may be contained in another signal-bearing media, such as a magnetic data storage diskette 700 (Figure 7), directly or indirectly accessible by the CPU 211.

10 Whether contained in the diskette 700, the computer/CPU 211, or elsewhere, the instructions may be stored on a variety of machine-readable data storage media, such as DASD storage (e.g., a conventional "hard drive" or a RAID array), magnetic tape, electronic read-only memory (e.g., ROM, EPROM, or EEPROM), an optical storage device (e.g. CD-ROM, WORM, DVD, digital optical tape, etc.), paper "punch" cards, or other suitable signal-bearing media
15 including transmission media such as digital and analog and communication links and wireless.

The second exemplary aspect of the present invention can be embodied in a number of variations, as will be obvious once the present invention is understood. That is, the methods of the present invention could be embodied as a
20 computerized tool stored on diskette 700 that contains a series of matrix subroutines to solve scientific and engineering problems using matrix processing.
YOR920030331US1

Alternatively, diskette 700 could contain a series of subroutines that allow an existing tool stored elsewhere (e.g., on a CD-ROM) to be modified to incorporate one or more of the features of the present invention which will be described shortly. It should also be obvious to one of skill in the art that the matrix subroutines described or referred to herein can be downloaded through a network interface from a remote storage facility.

The second exemplary aspect of the present invention additionally raises the issue of general implementation of the present invention in a variety of ways.

For example, it should be apparent, after having read the discussion above that the present invention could be implemented by custom designing a computer in accordance with the principles of the present invention. For example, an operating system could be implemented in which linear algebra processing is executed using the principles of the present invention.

In a variation, the present invention could be implemented by modifying standard matrix processing modules, such as described by LAPACK, so as to be based on the principles of the present invention. Along these lines, each manufacturer could customize their BLAS subroutines in accordance with these principles.

It should also be recognized that other variations are possible, such as versions in which a higher level software module interfaces with existing linear

YOR920030331US1

algebra processing modules, such as a BLAS or other LAPACK module, to incorporate the principles of the present invention.

Moreover, the principles and methods of the present invention could be embodied as a computerized tool stored on a memory device, such as independent
5 diskette 700, that contains a series of matrix subroutines to solve scientific and engineering problems using matrix processing, as modified by the technique described above. The modified matrix subroutines could be stored in memory as part of a math library, as is well known in the art. Alternatively, the computerized tool might contain a higher level software module to interact with existing linear
10 algebra processing modules.

It should also be obvious to one of skill in the art, after having read the present specification, that the instructions for the technique described herein can be downloaded through a network interface from a remote storage facility.

All of these various embodiments are intended as included in the present
15 invention, since the present invention should be appropriately viewed as a method to enhance the computation of matrix subroutines, as based upon recognizing how linear algebra processing can be more efficient by using the principles of the present invention.

In yet another exemplary aspect of the present invention, it should also be
20 apparent to one of skill in the art, after having read the present specification, that

YOR920030331US1

the principles of the present invention can be used in yet another environment in which parties indirectly take advantage of the present invention.

For example, it is understood that an end user desiring a solution of a scientific or engineering problem may undertake to directly use a computerized linear algebra processing method that incorporates the method of the present invention. Alternatively, the end user might desire that a second party provide the end user the desired solution to the problem by providing the results of a computerized linear algebra processing method that incorporates the method of the present invention. These results might be provided to the end user by a network transmission or even a hard copy printout of the results.

The present invention is intended to cover all these various methods of using the present invention, including the end user who uses the present invention indirectly by receiving the results of matrix processing done in accordance with the principles of the present invention.

The present invention provides a generalized technique that improves performance for linear algebra routines. The method and structure discussed here, yields higher-performance linear algebra routine processing for traditionally-slow subroutines that process matrices by providing a generalized data streaming technique that permits matrix data to efficiently be available for the matrix processing operation.

YOR920030331US1

All of the various embodiments discussed above are intended as included in the present invention, since the present invention should be appropriately viewed as methods to provide efficiency in the computation of conventional matrix subroutines by streaming data efficiently from higher levels of cache during execution of linear algebra subroutines.

While the invention has been described in terms of exemplary embodiments, those skilled in the art will recognize that the invention can be practiced with modification within the spirit and scope of the appended claims.

Further, it is noted that, Applicants' intent is to encompass equivalents of all claim elements, even if amended later during prosecution.